

516-38
237826

547682
PS

N90-12802

TDA Progress Report 42-98

August 15, 1989

68.

A Simple Model for DSS-14 Outage Times

H. C. Rumsey,¹ R. Stevens, and E. C. Posner

Office of Telecommunications and Data Acquisition

A model is proposed to describe DSS-14 outage times. Discrepancy Reporting System outage data for the period from January 1986 through September 1988 are used to estimate the parameters of the model. The model provides a probability distribution for the duration of outages, which agrees well with observed data. The model depends only on a small number of parameters, and has some heuristic justification. This shows that the Discrepancy Reporting System in the DSN can be used to estimate the probability of extended outages in spite of the discrepancy reports ending when the pass ends. The probability of an outage extending beyond the end of a pass is estimated as around 5 percent.

I. Introduction

A model is proposed to describe DSS-14 outage times. Outage data for the period from January 1986 through September 1988 are used to estimate the parameters of the model. The model provides a probability distribution for the duration of the outages. The model does not address questions about the mean time between outages. However, it does allow estimation of the probability of major outages even though the Discrepancy Reporting (DR) system stops recording them at the end of a pass. The philosophy has been that the best model is the simplest one that fits well enough.

The nature of the model, as above, is affected by the way outages are reported. If an extended outage occurs,

the time to restore service is not reported. Only the time lost for that pass is recorded. This limitation makes it difficult to determine the actual outage durations from the DR data and must be accounted for in the model. In the model, it is assumed that the actual time to restore service is hidden by a "cutoff" process which corresponds to the end of the pass. A typical pass runs for 9 hours. This means that the actual time to restore service is masked by a 9-hour cutoff window. (In developing the model, 8-, 9-, and 10-hour cutoffs were tried, with 9 fitting best and being reasonable on other considerations.)

II. Outage Distribution

Let $R(t)$ be the distribution function for *reported* outage durations, and let $A(t)$ be the distribution function for *actual* outage durations. Assuming the "cutoff" process and outage durations are independent, then for $t \geq 0$,

¹ Consultant to the TDA Office from Laguna Niguel, California.

$$\Pr(\text{reported time} \geq t)$$

$$= \Pr(\text{actual time} \geq t) \Pr(\text{cutoff time} \geq t)$$

or

$$1 - R(t) = (1 - A(t))(1 - t/U)_+ \quad (1)$$

where U is the duration of the pass (nominally 9 hours) and $(1 - t/U)_+$ is 0 for $t > U$. Equation (1) assumes that the start time of an outage is uniformly distributed over the duration of the pass.

Equation (1) deals with the “cutoff” problem but says nothing about the actual distribution of outages $A(t)$. The model for $A(t)$ is based on the reported outage data and a desire to minimize the number of parameters in the model. Figure 1 shows this measured outage data from the DR system, accurate to one minute. Thus, some minutes show multiple outages. There are 498 outages presented in this figure. The mean time to restore service is about 40 minutes.

The simplest model to fit such data would be an exponential distribution. The distribution function of an exponential random variable X of a mean a is given by

$$F(t) = \Pr(X \leq t) = 1 - e^{-t/a}$$

The fit is not accurate for short outages because Fig. 1 shows the density goes to 0 at 0-length outages, but an exponential has *maximum* density at 0-length outages. An exponential random variable with mean around 40 minutes fits the first part of the outage data fairly well beyond a few minutes up to about 100 minutes, but there are too many extreme values ($t \geq 100$ minutes) in the outage data. This suggests that there are two (or more) classes of failures, in addition to the short failures. For the first class of long failures, service can be restored quickly, in less than 40 minutes on the average. The second class of long failure requires more time to overcome.

III. Model

The following model has been adopted, with three parameter α , a , and b to be estimated, since U is 9 hours and not estimated:

$$\Pr(\text{reported time} \geq t) = 1 - R(t)$$

$$= ((1 - \alpha)(1 + t/a)e^{-t/a}$$

$$+ \alpha(1 + t/b)e^{-t/b})(1 - t/U)_+ \quad (2)$$

The form for this tail distribution has been taken to make the density function (essentially) zero at $t = 0$. Observe that if T_a is a random variable with distribution function $1 - (1 + t/a)e^{-t/a}$, then the expected value of T_a is $2a$ while the maximum of the density function of T_a occurs at $t = a$.

The parameters a and b occur symmetrically in Eq. (2). If a is chosen to be the smaller value, then outages of class “a” can be arranged to peak around $t = 10$ minutes. Outages of class “b” can be arranged to fit the tails of the observed outages. More heuristics appear in Section VI.

IV. Parameter Estimation

Equation (2) defines the model. To complete the model, good values for the parameters U , α , a , and b must be found. The maximum likelihood method will be used to estimate the parameters. Let (t_1, \dots, t_n) be the reported outage times. The t s are reported by the DR system to the nearest minute. Let t be one of the outage times. The models specify the probability $p(t)$ that an outage has duration t , where t is measured in minutes. Namely,

$$p(t) = R(t + 1/2) - R(t - 1/2)$$

where R is determined by Eq. (2). The probability that the observed outages occurred is the product of the probabilities of the separate outages,

$$\prod_{j=1}^n p(t_j)$$

This product is the likelihood function of the observations. It is a function of the model parameters U , α , a , and b . Maximum likelihood says to choose these parameters to maximize this product. The maximization is easy in this case because there are so few parameters.

There is a minor problem in determining the best value for the parameter U . Recall that U is the cutoff time for the pass. Finding U is like finding the end point of an interval $(0, U)$ in which a uniform random variable occurs. A little thought shows that the likelihood function for such a problem is maximized by taking U to be as small as possible. In the case of interest here, this would correspond to taking U to be slightly less than 8 hours (the maximum reported outage is 462 minutes). If extreme outages were common, this "defect" in the maximum-likelihood method would be no problem. Yet, the number of extended outages is small. So rather than estimate U from the likelihood function, U has been taken to be 9 hours throughout. To see how this selection affects the results, $U = 8$ and 10 hours for the case $F_b \equiv 1$ were also tried. As expected, $U = 8$ hours gives a larger value for the likelihood function, but the other parameters α and a are hardly changed. The results for $U = 10$ hours were not as good. Only $U = 9$ hours is considered below. This is consistent with the known distribution of pass lengths.

V. Goodness of Fit

A grid of points was used to find the maximum-likelihood values for the parameters. The maximum-likelihood values found were

$$\alpha = 0.186$$

$$a = 11.4 \text{ minutes}$$

$$b = 77.5 \text{ minutes}$$

The corresponding distribution function is then given by

$$\text{Pr}(\text{reported outage} \geq t) =$$

$$\left[0.814(1 + t/11.4)e^{-t/11.4} + 0.186(1 + t/77.5)e^{-t/77.5} \right] (1 - t/540) +$$

This equation represents the model. To see how well this model fits the observed outage data, the outage data was smoothed with a 5-point smoothing filter. The same filter was applied to the model as well. The results are dis-

played in Fig. 2. Qualitatively, the model fits the observed outages very well, for short, medium, and long outages.

The mean time MTR to restore service for the model is given by

$$\text{MTR} = (1 - \alpha)aT(U/a) + \alpha bT(U/b)$$

where

$$T(x) = 2 - 3/x + (1 + 3/x)e^{-x}$$

Substituting for α , a , b , and $U = 9$ hours (540 minutes) as always gives

$$\text{MTR} = 40.6 \text{ minutes}$$

This is in excellent agreement with the observed MTR of 40.4 minutes.

The probability that an outage exceeds 150 minutes was computed. With 498 outages, the model predicts there should be 28.4 outages of duration 150 minutes or longer. In the actual data of 498 outages there were 24 outages of this duration. Since the extended outage statistics are expected to be Poisson with an estimated mean of 28.4 and thus a sigma of $\sqrt{28.4} = 5.33$, the discrepancy of 4.4 is less than one sigma. This fits as well as could be expected with only 24 events. The probability of short and medium outages is also seen to fit very well. Hence, the use of the model seems indicated.

VI. Heuristics

It can be expected that the density of outages near zero outage time is very nearly 0, because it takes some minimum time to notice an outage and to respond to it, even by switching in a hot standby automatically. The $(1 + t/a)e^{-t/a}$ term in Eq. (2) does just this—it has density 0 at 0, for the corresponding density is $(t/a^2)e^{-t/a}$. This density also covers intermediate outages, but so does the "b" class. One might think that yet another distribution should be mixed in to cover these intermediate outages that cannot be recovered merely by switching something

in. But as has been seen, adding an extra one or two parameters is not necessary—the fit is very good with just the single time parameter a , the location of the maximum density of the short and medium outages, and the additional parameter α which gives the relative fraction of long outages.

The long outages are described by a similar distribution $(1 + t/b)e^{-t/b}$ with $b \gg a$. Why the tail of this and the shorter-outage distribution should be exponential is less clear, but the fit is good, and hard to tell from a distribution with a long low constant tail given the amount of data. It can be observed that the form of the a and b distribution arises as a difference of pure exponentials with infinitesimally close memoryless repair rates, but this does not seem to help the heuristics. The long tail can arise from certain failures such as low-noise maser warm up that takes a certain minimum time, e.g., 12 hours to recover from, when hot standbys for switching in are not provided.

Finally, as explained, the $(1 - t/U)_+$ multiplier term in Eq. (2) arises from the truncation of outage data at the end of a pass, where it was assumed that failures occur uniformly over the duration ($U = 9$ hours) of a pass. It is this truncation which makes it hard to distinguish a negative exponential tail from a long flat tail. The three-parameter model has been adopted even though the heuristics are not perfect.

VII. Summary

It has been shown that the Discrepancy Reporting System in the DSN can be used to give good estimates of the probability of extended outages in spite of the discrepancy reports ending when the pass ends. The probability of a major outage (one extending beyond the end of a pass) is estimated by the best-fit model as around 5 percent. The model also gives good estimates for the probability of short, medium, and long outages. It is simple and yet fits very well.

Acknowledgments

We thank D. W. Ginavan and G. C. Smith of the Bendix Field Engineering Corp. for organizing and providing the outage data used in this study.




Fig. 1. DSS-14 outage times, January 1986–September 1988.




Fig. 2. Smoothed DSS-14 outage times, January 1986–September 1988 compared with model.